

This application is submitted in the name of inventors Wei-Ping Sun, Beau James, and Eric Fronberg, assignors to Cisco Technology, Inc., a California Corporation.

5

S P E C I F I C A T I O N

TITLE OF INVENTION

METHOD AND APPARATUS FOR NETWORK MANAGEMENT SYSTEM

10

BACKGROUND OF THE INVENTION

Field of the Invention

The present invention relates to network management. More particularly, the present invention relates to a method and apparatus for providing a single point of control in a network management system.

15

Background

Network management systems are used to provide accessibility and security for local area networks (LANs). Such systems may include several components in a chassis or other configuration. The major components of a network management system include 20 system controllers and router cards, which are connected within the system by a bus. The router cards perform routing operations, including network security. System controllers coordinate the activities of the components of the system and provide a point of control for the system manager.

Ideally, a network management system should operate with a minimum of downtime and should be flexible enough to allow a system manager to modify the system while it is operating. Therefore, network management systems typically have redundant components which are capable of being replaced while the system is operating.

5 Redundant components may include system controllers and router cards.

Adding or replacing router cards in a system requires a method of providing suitable configurations and operating systems to a new router card so that the card can operate effectively in the system. Similarly, if a router card fails and a backup router card
10 must be pressed into service to replace the failed card, the backup card must be provided with a suitable configuration and operating system. Currently, such provision of configurations and operating systems is often performed by the system manager. Further, the manager may have to individually program each new router card if a single point of control is not available.

15

It is desirable that a network management system have a single point of control. With a single point of control, centralized in a system controller, a system manager can avoid having to program all system components separately. Further, it is desirable the single point of control provide the necessary files to the router cards automatically,
20 without need for intervention by the system manager.

BRIEF DESCRIPTION OF THE INVENTION

A network management system has a bus and a system controller connected to the bus. The bus has a plurality of fixed ports, where each of the ports is capable of connecting a router card to the system controller. A method of sending a file from the system controller to a router card includes receiving a request packet from the router card where the request packet includes a file type and the port address, identifying the file from the file type and the port address, and sending the file to the router card. The system controller has a memory storing the location of files associated with each of the port addresses, a request receiver for receiving a request for a file of a file type from a router card, an identifier for identifying a file, and a sender for sending the identified file to the router card. The router card has a memory for storing an operating system image and a configuration, and a request sender for sending a request for an operating system image and a configuration to the system controller upon power-up of the router card.

15

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram showing the functional components of the system in accordance with a specific embodiment of the present invention.

FIG. 2 is a block diagram showing the functional components of a system controller in
5 accordance with a specific embodiment of the present invention.

FIG. 3 is a block diagram showing the functional components of a router card in accordance with a specific embodiment of the present invention.

FIG. 4 is a process flow diagram showing the steps in a method of sending a file from the system controller to a router card in accordance with a specific embodiment of the

10 present invention.

FIG. 5A is diagram of a system packet according to a specific embodiment of the present invention.

FIG. 5B is diagram of a system frame header according to a specific embodiment of the present invention.

15 FIG. 6 is a process flow diagram showing the steps of establishing and continuing a connection between client and server in accordance with a specific embodiment of the present invention.

FIG. 7 is a process flow diagram showing the steps in the protocol in the event of a lost acknowledgement packet in accordance with a specific embodiment of the present
20 invention.

FIG. 8 is a process flow diagram showing the steps in the protocol in the event of a lost data packet in accordance with a specific embodiment of the present invention.

FIG. 9 is a diagram showing the format of a request packet in accordance with a specific embodiment of the present invention.

FIG. 10 is a diagram showing the format of a data packet in accordance with a specific embodiment of the present invention.

5 FIG. 11 is a diagram showing the format of a acknowledgement packet in accordance with a specific embodiment of the present invention.

FIG. 12 is a diagram showing the format of an error packet in accordance with a specific embodiment of the present invention.

DETAILED DESCRIPTION OF A PREFERRED EMBODIMENT

Those of ordinary skill in the art will realize that the following description of the present invention is illustrative only and not in any way limiting. Other embodiments of the invention will readily suggest themselves to such skilled persons having the benefit of 5 this disclosure.

The invention comprises a network management system and its associated components and methods. FIG. 1 is a block diagram showing the functional components of the system in accordance with an embodiment of the present invention. The system 10 includes a bus (10) and a system controller (12) connected to the bus. The bus has a plurality of fixed ports (14, 16, 18, 20, 22) capable of connecting a router card to the system controller via the bus. Two router cards (24, 26) are shown connected to the bus (10) at ports 14 and 16. The system may also include a backup system controller (28) connected to the bus, as well as other components connected to the bus. The system may 15 be organized in a chassis to facilitate changing of components. When the system is organized in a chassis, the bus may be incorporated into the back plane of the chassis.

The system includes a system controller having memory and processing capabilities. FIG. 2 is a block diagram showing the functional components of a system controller in accordance with an embodiment of the present invention. The system 20 controller (12) is connected to the system bus (10). The system controller (12) has a memory (30) storing the location of files associated with each of the port addresses on the bus (10). The files associated with the ports are of types including configuration files

and operating system files. The system controller includes a request receiver (32) for receiving a request for a file of a file type from a router card connected to the system controller at one of the ports. An identifier (34) communicates with the request receiver (32). The identifier (34) uses information in the memory (30), as well as information included in the request, to identify a file. A sender (36) sends the identified file to the router card which made the request.

In the system controller (12), the sender (36) includes an encapsulator (38) which encapsulates the file into a series of data packets. The sender (36) also includes a transmitter (40) for transmitting the data packets one at a time. An acknowledgment receiver (42) receives acknowledgments from the router card for individual data packets and causes the transmitter (40) to re-transmit a data packet if a proper acknowledgment for the data packet is not received. The acknowledgment receiver (42) causes the transmitter (40) to transmit the next data packet if a proper acknowledgment for a data packet is received.

The system controller (12) also includes a router card power controller (44) for instructing the router card to power down and then power up. A router card warm boot controller (46) is capable of instructing the router card to perform a warm boot. The system controller (12) may be designed so that either the router card power controller (44) or the router card warm boot controller (46) operates upon detecting the failure of another router card connected to the bus.

The system controller (12) may include memory for storing the requested files (48), or the files may be located at a remote file server (50). As described below, the request received by the request receiver (32) consists essentially of a system frame header, a request code, and a file type.

5

One or more router cards are connected to the bus in the system. Each router card has memory and processing capabilities. FIG. 3 is a block diagram showing the functional components of a router card in accordance with an embodiment of the present invention. The router card (24) is removably connected to the bus (10) at a port (14).

10 The router card (24) includes a memory (60) for storing files, including an operating system image and a configuration. The memory (60) is essentially empty upon power-up of the router card.

A request sender (62) sends requests for files including an operating system image
15 and a configuration to the system controller upon power-up or warm booting of the router card (24). When using the protocol described below, each request packet consists essentially of a system frame header, a request code, and a file type. A single request packet may be used for each file, or a single request packet may be used to request multiple files.

20

A data packet receiver (64) receives data packets comprising the requested files from the system controller, and a data packet checker (66) checks whether each data packet is received in good order. An acknowledgment sender (68) sends an

acknowledgment for a data packet to the system controller if the data packet received is in good order. The acknowledgement sender (68) re-sends a duplicate acknowledgment for the previous data packet to the system controller if a data package is not received.

When using the protocol described below, such acknowledgments consist essentially of

5 an system frame header, an acknowledgment code, and a block number.

The router card (24) may have a power control receiver (70) for receiving an instruction from the system controller to first power down and then power up.

Alternatively, the router card (24) may have a warm boot control receiver (72) for

10 receiving an instruction from the system controller to perform a warm boot.

The invention includes a method of transferring files from the system controller to a router card. FIG. 4 is a process flow diagram showing the steps in a method of sending a file from the system controller to a router card in accordance with a specific

15 embodiment of the present invention. The process may be triggered by powering a router card on (80) or warm booting the router card (82). Both of these events are significant in that the router card begins the process in a state where its memory for storing a configuration file and operating system are essentially empty. Typically, this situation arises when a new router card is installed in the system or when a router card is pressed into service for another router card which has gone out of service. When a new router card is installed, it may be powered on manually. When a router card is pressed into service as a backup, the system controller may direct it to power off and on or to perform a warm boot.

The file transfer process begins when a router sends a request to the system controller (84). The request is in the form of a request packet, as described below, and includes a file type. The request includes a system frame header and a request code,

5 where the system frame header include the address of the system controller and the port address of the router card and is less than or equal to 12 bytes in size.

The system controller receives the request packet from the router card (86). The system controller identifies the file or files (88) from the file type and the port address.

10 Once identified, the file is sent from the system controller to the router card (90).

When using the protocol described below, the process of sending the file includes several discrete steps. First, file parameters may be sent by the system controller (91) and received by the router card (92) after the file is identified. These parameters indicate

15 the file size and filename of the file to be transferred. Then, the system controller forms a data packet from the file (93), where the data packet is of a fixed size and includes a system frame header and a data packet protocol header. The data packet protocol header consists essentially of an operation code, a block number, a file type and a checksum.

The system frame header specifies the addresses of the router card and the system

20 controller and is less than or equal to 12 bytes in size. The data packet is transmitted by the system controller (94) to the router card.

If the data packet is lost, the router card will time out (95). The router card then sends a duplicate acknowledgement for the previous data packet (96) which is received by the system controller (97). This signals the system controller to transmit the data packet again (94). The time out occurs because the router does not receive a new data 5 packet within a certain time after the router card transmitted its last acknowledgement.

The router card receives the data packet (98). Then, the router card checks the content of the data packet (99). If there is an error in the content of the packet, the router card sends an error message (100) which is received by the system controller (101), 10 thereby causing the data packet to be transmitted again (94). If the router card determines that the data packet is in good order, the router card sends an acknowledgement for that packet (102). If the acknowledgement is lost, the system controller times out (103) because the system controller does not receive a new acknowledgement within a certain time after the system controller transmitted its last data packet. The time out causes the 15 system controller to re-transmit the data packet (94). When the acknowledgement is received by the system controller (104), the system controller sends the next data packet from the file, using the steps 93-104.

The last data packet in the file is formed having a size less than the fixed size of 20 the other data packets (105), signaling the router card that the file transfer is complete. The final packet is transmitted to the router card (106) using the same control mechanisms as for the other data packets, except as described below.

THE PROTOCOL

The protocol is used to transfer files between the system controller and the router card via the system interface. The interface is designed to allow communication between system components, including the file transfer operations described herein. The system interface transmits packets of fixed length between system components, and system components using the interface have system buffers of a size equal to the fixed length.

10 The protocol is designed to be small and easy to implement. Therefore, it lacks some of the features of a regular file transfer protocol. For example, it has no provisions for user authentication. The only thing it can do is to mail packet to and from between the system controller and the router card.

15 Overview of the Protocol

The device that begins with a request is the router card, or client, and the device that grants the request is the system controller, or server. If the server grants the request from the client, the connection is opened and the file is sent in fixed length blocks of the 20 length of the system buffer size. Each data packet contains one block of data, and must be acknowledged by a proper acknowledgment packet before the next packet can be sent. To be proper, an acknowledgement packet must be received by the system controller and include the proper block number, source, and destination as described below. A data

packet of less than the length of the system buffer size signals termination of a transfer. If a packet gets lost in the network, the intended recipient will timeout and may retransmit his last packet (which may be data or an acknowledgment), thus causing the sender of the lost packet to retransmit that lost packet. The sender has to keep just one packet on hand
5 for retransmission, since the lock step acknowledgment guarantees that all older packets have been received. Notice that both nodes involved in a transfer are considered senders and receivers. One sends data and receives acknowledgments, the other sends acknowledgments and receives data.

10 Most errors cause termination of the connection. An error is signaled by sending an error packet. This packet is not acknowledged, and not retransmitted, so the other end of the connection may not get it. Therefore timeouts are used to detect such a termination when the error packet has been lost. Errors are caused by three types of events: not being able to satisfy the request (e.g., file not found, access violation), receiving a packet which
15 cannot be explained by a delay or duplication in the network (e.g., an incorrectly formed packet), and losing access to a necessary resource (e.g., memory full or access denied during a transfer).

The protocol is very restrictive, in order to simplify implementation. For example,
20 the fixed length blocks make allocation straightforward, and the lock step acknowledgement provides flow control and eliminates the need to reorder incoming data packets.

The protocol is designed to be implemented on top of the system interface. The data packets will have a system frame header and a protocol header. FIG. 5A is diagram of a system packet according to a specific embodiment of the present invention. As shown in FIG. 5A, the order of the contents of a packet will be: system frame header
5 (110), protocol header (112), followed by the remainder of the packet (114). FIG. 5B is diagram of a system frame header according to a specific embodiment of the present invention. The source (113) and destination (115) port numbers are specified in the system frame header.

10 The protocol header consists of a 2 byte opcode field which indicates the packet's type (e.g., DATA, ERROR, etc.) These opcodes and the formats of the various types of packets are discussed further in the section on protocol packets.

Initial Connection

15 FIG. 6 is a process flow diagram showing the steps of establishing and continuing a connection between client and server in accordance with an embodiment of the present invention. Note that RRQ, ACK, and DATA are the names of the read request, acknowledgment, and data types of packets respectively.

20 A transfer is established by sending a request from the client to the server (120). The request is received by the server (122). A positive response from the server is a special data packet called DATA[0] (124). By way of background, each data packet has associated with it a block number. Block numbers are consecutive and begin with one. In

the special data packet the block number is zero. The special data packet contains file parameters, namely, the filename and file size as well as a checksum.

The client responds to data packets with acknowledgments. An acknowledgment
5 packet will contain the block number of the data packet being acknowledged. After receiving the special data packet (126), the client responds to the special data packet with an acknowledgement packet indicating block number zero (128). The server receives this acknowledgement (130). Subsequently, the server sends data packets beginning with block one (132) and, if all goes well, the data packets are received by the client (134) and
10 acknowledgements are sent by the client (136).

During the file transfer process, system port addresses are agreed upon as follows.

The router card sends a "RRQ" to the system controller with destination in the system
15 frame header equal to the system controller port number. The source in the system frame header will identify the port address of the router card. The system controller then sends a "DATA" (with block number = 0) to the router card with source in the system frame header equal to the router card port number. The router card sends a "ACK" (with block number = 0) to the system controller.

20

After the connection has been established, the first data packet can be sent by the system controller with a block number of 1. In the next step, and in all succeeding steps, the system controller should make sure that the source port number matches the value

that was agreed on. If the source port number does not match, the packet should be discarded as erroneously sent from somewhere else. An error packet should be sent to the source of the incorrect packet, while not disturbing the transfer.

5 FIG. 7 is a process flow diagram showing the steps in the protocol in the event of a lost acknowledgement packet in accordance with a specific embodiment of the present invention. If an acknowledgement packet is lost (140), the server will time out (142) and retransmit the last data packet it sent (144). The client receives the duplicate data packet (146), and sends a duplicate of the lost acknowledgement (148). Upon receiving the
10 proper acknowledgement (150), the server transmits the next data packet (152).

FIG. 8 is a process flow diagram showing the steps in the protocol in the event of a lost data packet in accordance with a specific embodiment of the present invention. If a data packet is lost (160), the client will time out (162) and retransmit the last
15 acknowledgement packet it sent (164). The server receives the duplicate acknowledgement packet (166), and sends a duplicate of the lost data packet (168). The duplicate data packet is received by the client (170) and an acknowledgement is sent (172).

20 The Protocol Packets

The protocol supports four types of packets, all of which have been mentioned

above. The protocol header of a packet contains the opcode associated with that packet, as shown in Table 1.

opcode	operation
5	1 Read request (RRQ)
	2 Data (DATA)
	3 Acknowledgment (ACK)
	4 Error (ERROR)

Table 1

FIG. 9 is a diagram showing the format of a request packet in accordance with a specific embodiment of the present invention. The request packet includes a system frame header (180), the RRQ opcode (182), and a file type (184). The type is the requesting file type, for example TYPE_IMAGE for operating system image file or TYPE_CONFIG for configuration file. The type is user defined and is recognized by both client and server. If the type is not recognized or is not allowed by the receiver then the packet should be discarded and an error packet should be sent to the source of the incorrect packet. The protocol allows only one connection between client and server at any one time. The second connection request should be rejected if the first connection is still alive. The first connection will be maintained while the second is rejected by returning an error packet.

Data is transferred in DATA packets. FIG. 10 is a diagram showing the format of a data packet in accordance with a specific embodiment of the present invention. The data packet includes a system frame header (190), the DATA opcode (192), the data block number (194), the file type (196), a checksum (198), and data (200). The opcode, 5 block number, type, and checksum may collectively be referred to a data packet protocol header (202). The block numbers on data packets begin with one and increase by one for each new block of data. The block number on data packet with zero is a special packet responses to the RRQ. The data packet with block number zero contains parameters for the client to initialize itself properly. The parameters are defined for different types and 10 need to be recognized by both client and server. For example, the parameters for TYPE_IMAGE can be filename, file size and file checksum. The checksum is the checksum number of each data packet. The receiver generates a checksum on the data packet after receiving the packet and compares it to the checksum number in the checksum field. If the number is different then the packet is dropped and error packet is 15 sent to server to request retransmission. The data field is from zero byte to system buffer size - headers (system header + data packet protocol header) size long. If it is equal to system buffer size - header size bytes, the block is not the last block of data; if it is from zero bytes to data field size -1 long, it signals the end of the transfer.

20 All DATA packets are acknowledged by ACK or ERROR packets unless a timeout occurs. The RRQ and ACK packets are acknowledged by DATA or ERROR packets. FIG. 11 is a diagram showing the format of an acknowledgement packet in accordance with a specific embodiment of the present invention. The acknowledgement

packet includes a system frame header (210), the acknowledgement packet opcode (212), and the data block number (214). The block number in an ACK echoes the block number of the DATA packet being acknowledged.

5 FIG. 12 is a diagram showing the format of an error packet in accordance with a specific embodiment of the present invention. The error packet includes a system frame header (220), the error opcode (222), a code describing the specific type of error (224), an error message (226), and a zero byte (228). An ERROR packet can be the acknowledgment of any other type of packet. The error code is an integer indicating the
10 nature of the error. The error message is intended for human consumption, and should be in netascii. It is terminated with a zero byte. Both the system controller and the router card have error message senders and error message receivers to implement the sending and receiving of error messages.

15 Normal Termination .

The end of a transfer is marked by a DATA packet that contains between 0 and data field size minus one of data. This packet is acknowledged by an ACK packet like all other DATA packets. The server acknowledging the final DATA packet may terminate
20 its side of the connection on sending the final ACK. On the other hand, delaying is encouraged. This means that the server sending the final ACK will wait for a while before terminating in order to retransmit the final ACK if it has been lost. The acknowledger will know that the ACK has been lost if it receives the final DATA packet

again. The server sending the last DATA must retransmit it until the packet is acknowledged or the sending server times out. If the response is an ACK, the transmission was completed successfully. If the sender of the data times out and is not prepared to retransmit any more, the transfer may still have been completed successfully,

5 after which the acknowledger or network may have experienced a problem. It is also possible in this case that the transfer was unsuccessful. In any case, the connection has been closed.

Premature Termination

If a request can not be granted, or some error occurs during the transfer, then an ERROR packet (opcode 4) is sent. This is only a courtesy since it will not be retransmitted or acknowledged, so it may never be received. Timeouts must also be used to detect errors.

SYSTEM DESIGN

The system controller is designed to simultaneously handle multiple image download requests from router cards. The image download request will initialize a system controller image download process to handle the file transfer. It is a point to point connection and only one connection is allowed at any one time for each port address. The process will be terminated once the image download process completed. Multiple image download processes may be launched to handle multiple router card image requests.

System flow

The following summarizes the steps of image download.

- (a) The router card initiates the image download request to the system controller.
- (b) The system controller verifies the request and checks its memory for router card port number, image filename and router card root path. If the image can not be found or opened then the system controller responds with an error packet. Otherwise, the system controller responds with the image file size and name.
- (c) The router card allocates buffer to store the image and acknowledges with "ready to receive" message (ACK[0]).
- 10 (d) The system controller opens the image file, copies the first block (block size = system buffer size - header size) and sends to the router card.
- (e) The router card copies the data block to DRAM and acknowledges with data block number and checksum.
- 15 (f) The system controller verifies the checksum number and retransmits the message if there is a checksum error, otherwise copies the next block and sends to the router card.
- (g) Steps (e) and (f) are repeated until all messages are received by the router card.
- 20 (h) The system controller sends an end-of-file message with image checksum number to the router card for connection termination.
- (i) The router card verifies the total bytes received and system image checksum number. If it doesn't match then the router card re-requests an image.
- (j) The router card launches the operating system image.

A timer will be associated with each message for error (timeout) correction.

If a non-active router card (router card two) is designated as a backup to an active router card (router card one), the non-active router will run the same image of the active router when it becomes active. The active and non-active router cards may run different images initially. However, after failure of the router card one, router card two will be reloaded and request its image from the system controller. The system controller image download process will verify the redundancy information first when receiving the request, and find the router card one image and configuration and sends the files to

10 router card two. For example,

(a) Router card one at port 1 and router card two and port 2 run different images initially.

(b) The router cards are configured as a redundant pair with router card one as the active router card.

15 (c) Router card one crashes. The system controller recognizes the crash and causes router card two to power off and on or to warm boot.

(d) Router card two reloads and requests its image from the system controller.

(e) The system controller image download process verifies that port 1 and port 2 are a redundant pair and loads port 1 image configured in memory. Router card

20 two downloads the port 1 router card image and launches it.

It will be apparent that the functional components of the system controller and the router card may be implemented in a variety of ways. For example, they may be implemented in the form of program storage devices readable by a machine and tangibly embodying a program of instructions executable by the machine to perform the methods

5 described herein. Such implementations may include a variety of operating systems, computing platforms, and/or computer languages. In addition, those of ordinary skill in the art will readily recognize that devices such as hardwired devices, devices relying on FPGA (field programmable gate array) or ASIC (application specific integrated circuit) technology, or the like, may also be used without departing from the scope and spirit of

10 the inventive concepts described herein.

While embodiments and applications of this invention have been shown and described, it would be apparent to those skilled in the art having the benefit of this disclosure that many more modifications than mentioned above are possible without departing from the inventive concepts herein after a review of this disclosure. The

15 invention, therefore, is not to be restricted except in the spirit of the appended claims.